

---

*Research Article: Open Source Tools and Methods | Novel Tools and Methods*

## **PsychRNN: An Accessible and Flexible Python Package for Training Recurrent Neural Network Models on Cognitive Tasks**

<https://doi.org/10.1523/ENEURO.0427-20.2020>

**Cite as:** eNeuro 2020; 10.1523/ENEURO.0427-20.2020

Received: 21 October 2020

Revised: 22 November 2020

Accepted: 24 November 2020

---

*This Early Release article has been peer-reviewed and accepted, but has not been through the composition and copyediting processes. The final version may differ slightly in style or formatting and will contain links to any extended data.*

**Alerts:** Sign up at [www.eneuro.org/alerts](http://www.eneuro.org/alerts) to receive customized email alerts when the fully formatted version of this article is published.

Copyright © 2020 Ehrlich et al.

This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International license, which permits unrestricted use, distribution and reproduction in any medium provided that the original work is properly attributed.

1       PsychRNN: An Accessible and Flexible Python Package for  
2       Training Recurrent Neural Network Models on Cognitive Tasks

3

4       Daniel B. Ehrlich<sup>1,\*</sup>, Jasmine T. Stone<sup>2,\*</sup>, David Brandfonbrener<sup>2,3</sup>, Alexander Atanasov<sup>4,5</sup>,  
5       John D. Murray<sup>1,4,6,†</sup>

6       <sup>1</sup> *Interdepartmental Neuroscience Program, Yale University, New Haven, CT, USA*

7       <sup>2</sup> *Department of Computer Science, Yale University, New Haven, CT, USA*

8       <sup>3</sup> *Department of Computer Science, New York University, New York, NY, USA*

9       <sup>4</sup> *Department of Physics, Yale University, New Haven, CT, USA*

10      <sup>5</sup> *Department of Physics, Harvard University, Cambridge, MA, USA*

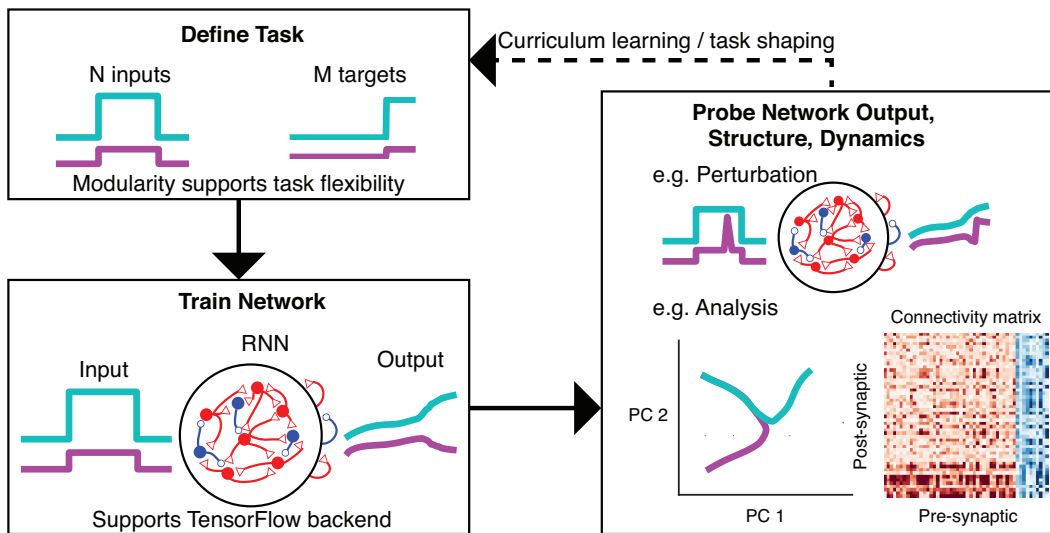
11      <sup>5</sup> *Department of Psychiatry, Yale School of Medicine, New Haven, CT, USA*

12      \* *Equal contribution*

13      † *Correspondence: john.murray@yale.edu*

14 **Abstract**

15 Task-trained artificial recurrent neural networks (RNNs) provide a computational modeling frame-  
16 work of increasing interest and application in computational, systems, and cognitive neuroscience.  
17 RNNs can be trained, using deep learning methods, to perform cognitive tasks used in animal  
18 and human experiments, and can be studied to investigate potential neural representations and  
19 circuit mechanisms underlying cognitive computations and behavior. Widespread application of  
20 these approaches within neuroscience has been limited by technical barriers in use of deep learning  
21 software packages to train network models. Here we introduce PsychRNN, an accessible, flexible,  
22 and extensible Python package for training RNNs on cognitive tasks. Our package is designed for  
23 accessibility, for researchers to define tasks and train RNN models using only Python and NumPy  
24 without requiring knowledge of deep learning software. The training backend is based on Ten-  
25 sorFlow and is readily extensible for researchers with TensorFlow knowledge to develop projects  
26 with additional customization. PsychRNN implements a number of specialized features to support  
27 applications in systems and cognitive neuroscience. Users can impose neurobiologically relevant  
28 constraints on synaptic connectivity patterns. Furthermore, specification of cognitive tasks has a  
29 modular structure, which facilitates parametric variation of task demands to examine their impact  
30 on model solutions. PsychRNN also enables task shaping during training, or curriculum learning, in  
31 which tasks are adjusted in closed-loop based on performance. Shaping is ubiquitous in training of  
32 animals in cognitive tasks, and PsychRNN allows investigation of how shaping trajectories impact  
33 learning and model solutions. Overall, the PsychRNN framework facilitates application of trained  
34 RNNs in neuroscience research.



**Visual Abstract:** Example workflow for using PsychRNN. First, the task of interest is defined, and a recurrent neural network model is trained to perform the task, optionally with neurobiologically informed constraints on the network. After the network is trained, the researchers can investigate network properties including the synaptic connectivity patterns and the dynamics of neural population activity during task execution, and other studies, e.g. those on perturbations, can be explored. The dotted line shows the possible repetition of this cycle with one network, which allows investigation of training effects of task shaping, or curriculum learning, for closed-loop training of the network on a progression of tasks.

### 35 **Significance Statement**

36 Artificial recurrent neural network (RNN) modeling is of increasing interest within computational,  
37 systems, and cognitive neuroscience, yet its proliferation as a computational tool within the field  
38 has been limited due to technical barriers in use of specialized deep-learning software. PsychRNN  
39 provides an accessible, flexible, and powerful framework for training RNN models on cognitive  
40 tasks. Users can define tasks and train models using the Python-based interface which enables  
41 RNN modeling studies without requiring user knowledge of deep learning software or comprehensive  
42 understanding of RNN training. PsychRNN's modular structure facilitates task specification and  
43 incorporation of neurobiological constraints, and supports extensibility for users with deep learning  
44 expertise. PsychRNN's framework for RNN modeling will increase accessibility and reproducibility  
45 of this approach across neuroscience subfields.

## 46 Introduction

47 Studying artificial neural networks (ANNs) as models of brain function is an approach of increasing  
48 interest in computational, systems, and cognitive neuroscience (Kriegeskorte, 2015; Yamins and DiCarlo,  
49 2016; Richards et al., 2019). ANNs comprise many simple units, called neurons, whose synaptic  
50 connectivity patterns are iteratively updated via deep learning methods to optimize an objective.  
51 For application in neuroscience and psychology, ANNs can be trained to perform a cognitive task  
52 of interest, and the trained networks can then be analyzed and compared to experimental data  
53 in a number of ways, including their behavioral responses, neural activity patterns, and synaptic  
54 connectivity. Recurrent neural networks (RNNs) form a class of ANN models which are especially  
55 well-suited to perform cognitive tasks which unfold across time, common in psychology and neuro-  
56 science, such as decision-making or working-memory tasks (Sussillo, 2014; Song et al., 2016; Barak,  
57 2017; Yang and Wang, 2020). In RNNs, highly recurrent synaptic connectivity is optimized to gener-  
58 ate target outputs through the network population dynamics. RNNs have been applied to model  
59 the dynamics of neuronal populations in cortex during cognitive, perceptual, and motor tasks and  
60 are able to capture associated neural response dynamics (e.g., Mante et al., 2013; Sussillo et al.,  
61 2015; Carnevale et al., 2015; Rajan et al., 2016; Remington et al., 2018; Masse et al., 2019).

62 Despite growing impact of RNN modeling in neuroscience, wider adoption by the field is cur-  
63 rently hindered by the requisite knowledge of specialized deep learning platforms, such as Tensor-  
64 Flow or PyTorch, to train RNN models. This creates accessibility barriers for researchers to apply  
65 RNN modeling to their neuroscientific questions of interest. It can be especially challenging in  
66 these platforms to implement neurobiologically motivated constraints, such as structured synaptic  
67 connectivity or Dale’s principle defining excitatory and inhibitory neurons. There is also need for  
68 modular frameworks to define the cognitive tasks on which RNNs are trained, which would facil-  
69 itate investigation of how task demands shape network solutions. To better model experimental  
70 paradigms for training animals on cognitive tasks, an RNN framework should enable investigation  
71 of task shaping, in which training procedures are progressive adapted to the subject’s performance  
72 during training.

73 To address these challenges, we developed the software package PsychRNN as an accessible,  
74 flexible, and extensible computational framework for training RNNs on cognitive tasks. Users  
75 define tasks and train RNN models using only Python and NumPy, without requiring comprehensive  
76 understanding of ANNs. The training backend is based on TensorFlow and is extensible for projects  
77 with additional customization. PsychRNN implements a number of specialized features to support  
78 applications in systems and cognitive neuroscience, including neurobiologically relevant constraints  
79 on synaptic connectivity patterns. Specification of cognitive tasks has a modular structure, which  
80 aids parametric variation of task demands to examine their impact on model solutions and promotes  
81 code reuse and reproducibility. Modularity also enables implementation of curriculum learning, or  
82 task shaping, in which tasks are adjusted in closed-loop based on performance. Our overall goal  
83 for PsychRNN is to facilitate application of RNN modeling in neuroscience research.

## 84 **Materials and Methods**

### 85 **Package structure**

86 To serve our objectives of accessibility, extensibility, and reproducibility, we divided the PsychRNN  
87 package into two main components: the `Task` object and the `Backend` (**Fig. 1**). We anticipate that  
88 all PsychRNN users will want to be able to define novel tasks specific to their research domains  
89 and questions. The `Task` object is therefore fully accessible to users without any TensorFlow or  
90 deep learning background. Users familiar with Python and NumPy are able to fully customize  
91 novel tasks, and they can customize network structure (e.g., number of units, form of nonlinearity,  
92 connectivity) through preset options built into the `Backend`.

93 For users with greater need for flexibility in network design, the `Backend` is designed for acces-  
94 sibility, customizability, and extensibility. `Backend` customization typically requires knowledge of  
95 TensorFlow. For those with TensorFlow knowledge, PsychRNN’s modular design enables definition  
96 of new models, regularizations, loss functions, and initializations. This modularity facilitates test-  
97 ing hypotheses regarding the impact of specific potential structural constraints on RNN training  
98 without having to expend time and resources designing a full RNN codebase.

### 99 **Task object**

100 The `Task` object is structured to allow users to define their own new task using Python and NumPy.  
101 Specifically, `generate_trial_params` creates trial specific parameters for the task (e.g., stimulus  
102 and correct response). `trial_function` specifies the input, target output, and output mask at  
103 a given time `t`, given the parameters generated by `generate_trial_params`. PsychRNN comes  
104 set with three example tasks that are well researched by cognitive neuroscientists: perceptual dis-  
105 crimination (Roitman and Shadlen, 2002), delayed discrimination (Romo et al., 1999), and delayed  
106 match-to-category (Freedman and Assad, 2006). These tasks highlight possible schemas users can  
107 apply to specifying their own tasks and provide tasks with which users can test the effect of different  
108 structural network features.

109 Tasks can optionally include `accuracy` functions. Accuracy measures performance in a manner  
110 more relevant to experiments than traditional machine learning measures such as loss. On a given  
111 trial, accuracy is either one (success) or zero (failure). In contrast, loss on a given trial is a real-  
112 numbered value. Accuracy is calculated over multiple trials to obtain a ratio of trials correct to  
113 total trials. Accuracy is used as the default metric by the `Curriculum` class.

### 114 **Backend**

115 The `Backend` includes all of the neural network training and specification details (**Fig. 1, Step 2**).  
116 The backend, while being accessible and customizable, was designed with pre-set defaults sufficient  
117 to get started with PsychRNN. The TensorFlow details are abstracted away by the `Backend` so  
118 that researchers are free to work with or without an understanding of TensorFlow. Additionally,  
119 since the `Backend` is internally modular, different components of the `Backend` can be swapped in  
120 and out interchangeably. In this section, modular components of the `Backend` are described so that

121 researchers who want to get more in-depth with PsychRNN know what tools are available to them.

## 122 **Models**

RNNs are a large class of ANNs that process input over time. In the PsychRNN release, we include a basic RNN (which we refer to as an RNN throughout the rest of the paper), and an LSTM model (Hochreiter and Schmidhuber, 1997). The basic RNN model is governed by the following equations:

$$\begin{aligned}\tau dx &= (-x + W_{rec}r + b_{rec} + W_{in}u)dt + \sigma_{rec}\sqrt{2\tau}d\xi \\ r &= f(x) \\ z &= W_{out}r + b_{out}\end{aligned}$$

123 where  $u$ ,  $x$  and  $z$  are the input, recurrent state, and output vectors, respectively.  $W_{in}$ ,  $W_{rec}$  and  
124  $W_{out}$  are the input, recurrent, and output synaptic weight matrices.  $b_{rec}$  and  $b_{out}$  are constant  
125 biases into the recurrent and output units.  $dt$  is the simulation time-step and  $\tau$  is the intrinsic  
126 timescale of recurrent units.  $\sigma_{rec}$  is a constant to scale recurrent unit noise, and  $d\xi$  is a gaussian  
127 noise process with mean 0 and standard deviation 1.  $f(x)$  is a nonlinear transfer function, which by  
128 default in PsychRNN is rectified linear (ReLU). This default can be replaced with any TensorFlow  
129 transfer function.

130 PsychRNN also includes an implementation of LSTMs (Long Short Term Memory networks), a  
131 special class of RNNs that enables longer-term memory than is easily attainable with basic RNNs  
132 (Hochreiter and Schmidhuber, 1997). LSTMs use a separate “cell state” to store information gated  
133 by sigmoidal units. Additional models can be user-defined but require knowledge of TensorFlow.

## 134 **Initializations**

135 The synaptic weights that define an ANN are typically initialized randomly. However, with RNNs,  
136 large differences in performance, training time and total asymptotic loss, have been observed for  
137 different initializations (V. Le et al., 2015). Since initializations can be crucial for training, we have  
138 included several initializations currently used in the field (Glorot and Bengio, 2010). By default,  
139 recurrent weights are initialized randomly from a gaussian distribution with spectral radius of 1.1  
140 (Sussillo and Abbott, 2009). We also include an initialization called Alpha Identity that initializes  
141 the recurrent weights as an identity matrix scaled by a parameter  $\alpha$  (V. Le et al., 2015). Each of  
142 these initializations can substantially improve the training process of RNNs. PsychRNN includes a  
143 `WeightInitialization` class that initializes all network weights randomly, all biases as zero, and  
144 connectivity masks as all-to-all. New initializations inherit this class and can override any variety  
145 of initializations defined in the base class.

## 146 **Loss functions**

147 During training an RNN is optimized to minimize the loss, so the choice of loss function can  
148 be crucial for determining exactly what the network learns. By default, the loss function is  
149 `mean_squared_error`. Our Backend also includes an option for using `binary_cross_entropy` as



150 the loss function. Other loss functions can be easily defined with some TensorFlow knowledge and  
151 added to the `LossFunction` class. Loss functions take in the network output (`predictions`), the  
152 target output (`y`) and the output mask, and return a float calculated using the TensorFlow graph.

### 153 **Regularizers**

154 Regularizers are penalties added to the loss function that may help prevent the network from  
155 overfitting to the training data. We include options for L2-norm and L1-norm regularization for  
156 the synaptic weights, which tend to reduce the magnitude of weights and sparsify the resulting  
157 weight matrices. In addition we include L2-norm regularization on the post-nonlinearity recurrent  
158 unit activity  $r$ . Other regularizations can be added to the `Regularizer` class through TensorFlow.  
159 By default, no regularizations are used.

### 160 **Optimizers**

161 PsychRNN is built to take advantage of the many optimizers available in the TensorFlow package.  
162 Instead of explicitly defining equations for back propagation through time (BPTT), PsychRNN  
163 converts the user supplied Task and RNN into a "graph" model interpretable by TensorFlow.  
164 TensorFlow can then automatically generate gradients of the user supplied `LossFunction` with  
165 respect to the weights of the network. These gradients can then be used by any TensorFlow  
166 optimization algorithm such as stochastic gradient descent, Adam or RMSProp to update the  
167 weights and improve task performance (Ruder, 2017).

### 168 **Neurobiologically motivated connectivity constraints**

169 PsychRNN is designed for investigation of neurobiologically motivated constraints on the input,  
170 recurrent, and output synaptic connectivity patterns. The user can specify which synaptic connec-  
171 tions are allowed and which are forbidden (set to zero) through optional user-defined masks at the  
172 point of RNN model initialization. This feature enables modeling neural architectures including  
173 sparse connectivity and multi-region networks (Rikhye et al., 2018). Optional user-defined masks  
174 allow specification of which connections are fixed in their weight values, and which connections are  
175 plastic for optimization during training (Rajan et al., 2016). By default, all weights are allowed to  
176 be updated by training. PsychRNN also enables implementation of Dale's principle, such that each  
177 recurrent unit's synaptic weights are all of the same sign (i.e., each neuron's post-synaptic weights  
178 are either all excitatory or all inhibitory) (Song et al., 2016). The optional parameter `dales_ratio`  
179 sets the proportion of excitatory units, making the balance inhibitory.

### 180 **Curriculum learning**

181 Curriculum learning refers to the presentation of training examples structured into successive dis-  
182 crete blocks sorted by increasing difficulty (Bengio et al., 2009; Krueger and Dayan, 2009). Task  
183 modularity in PsychRNN enables an intuitive framework for curriculum learning that does not  
184 require TensorFlow knowledge. Curriculum learning is implemented by passing a `Curriculum` ob-  
185 ject to the RNN model when training is executed. Although very flexible and customizable, the

186 simplest form of the `Curriculum` object can be instantiated solely with the list of tasks that one  
187 wants to train on sequentially.

188 The `Curriculum` class included in PsychRNN is flexible and extensible. By default, accuracy,  
189 as defined within a task, is used to measure the performance of the task. When the performance  
190 surpasses a user-defined threshold, the network starts training with the next task. The `Curriculum`  
191 object thus includes an optional input array, `thresholds`, for specifying the performance thresholds  
192 required to advance to each successive task. Apart from accuracy one may wish to advance the  
193 curriculum stage using an alternative measure such as loss or number of iterations. We include an  
194 optional `metric` function that can be passed into the `Curriculum` class to define a custom measure  
195 to govern task stage transitions.

## 196 Simulator

197 One limitation of specifying RNN networks in the TensorFlow language is that in order to run a  
198 network the inputs, outputs, and computation need to take place within the TensorFlow framework,  
199 which can impede users' ability to design and implement experiments on their trained RNN models.  
200 To mitigate this, we have included a NumPy-based simulator which takes in RNN and `Task` objects  
201 and simulates the network in NumPy. This simulator allows the user to study various neuroscientific  
202 topics such as robustness to perturbations.

## 203 Software availability

204 The PsychRNN open-source software described in the paper is available online for download in a  
205 Git-based repository at <https://github.com/murraylab/PsychRNN>. Detailed documentation con-  
206 taining tutorials and examples is also provided. The code and documentation are available as  
207 Extended Data. All data and figures included were produced on a MacBook Pro (Retina, 13-inch,  
208 Early 2015) with 8 GB of RAM and 2.7 GHz running macOS Catalina 10.15.5 in an Anaconda  
209 environment with Python 3.6.9, NumPy 1.17.2, and TensorFlow 1.14.0.

## 210 Results

211 To facilitate accessibility, PsychRNN allows users to define tasks and define and train networks using  
212 a Python- and NumPy-based interface. PsychRNN provides a machine-learning backend, based on  
213 TensorFlow, which converts task and network specifications into the Tensorflow deep learning frame-  
214 work to optimize network weights. This allows users to focus on the neuroscientific questions rather  
215 than implementation details of deep learning software packages. As an example, we demonstrate  
216 how PsychRNN can specify an RNN model, train it to perform a task of neuroscientific interest—  
217 here, a two-alternative forced-choice perceptual discrimination task (Roitman and Shadlen, 2002)—  
218 and return behavioral readout from output units and internal activity patterns of recurrent units  
219 (**Fig. 2**).

## 220 Modularity

221 The PsychRNN Backend is complimented by the **Task** object which enables users to easily and  
222 flexibly specify tasks of interest without any prerequisite knowledge of TensorFlow or machine  
223 learning. The **Task** object allows flexible input and output structure, with tasks varying in not only  
224 the task-specific features but also the number of input and output channels (**Fig. 3**). Furthermore,  
225 the object-oriented structure of task definition in PsychRNN facilitates tasks that can be quickly  
226 and easily varied along multiple dimensions. For example in an implementation of a delayed  
227 discrimination task (Romo et al., 1999), we can vary stimuli and delay durations with a set of two  
228 parameters (**Fig. 3B**). Importantly not only can we vary the inputs as they exist, but integration  
229 between the **Task** object and Backend makes it possible to vary the structure of the network from  
230 the **Task** object. This is due to the fact that the RNN models are constructed after task definitions  
231 using task parameters, and are therefore custom-structured to accommodate task features. In our  
232 implementation of a delayed match-to-category task (Freedman and Assad, 2006), we can freely  
233 change the number of inputs (input discretization) and the number of outputs (categories) (**Fig.**  
234 **3D**). This flexibility allows researchers to investigate how the network solution of trained RNNs  
235 may depend on task or structural properties (Orhan and Ma, 2019).

## 236 Neurobiologically motivated connectivity constraints

237 While there are multiple general-purpose frameworks for training ANNs, neuroscientific model-  
238 ing often requires neurobiologically motivated constraints and processes which are not common in  
239 general-purpose ANN software. PsychRNN includes a variety of easily implemented forms of con-  
240 straints on synaptic connectivity. The default RNN network has all-to-all connectivity, and allows  
241 units to have both excitatory and inhibitory connections. Users can specify which potential synaptic  
242 connections are forbidden or allowed, as well as which are fixed and which are plastic for updating  
243 during training. Furthermore, PsychRNN can enforce Dale’s principle, so that each unit has either  
244 all-excitatory or all-inhibitory synapses onto its targets. **Fig. 4** demonstrates that networks can  
245 be trained while subject to various constraints on recurrent connectivity. For example, units can  
246 be prevented from making autapses (i.e., self-connections). Networks with block-like connectivity  
247 matrices can be used to model multiple brain regions, with denser within-region connectivity and  
248 sparser between-region connectivity.

## 249 Curriculum learning

250 One important feature included in PsychRNN is a native implementation of curriculum learning.  
251 Curriculum learning, also referred to as task shaping in the psychological literature (Krueger and Dayan,  
252 2009), refers to structuring training examples such that the agent learns easier trials or more basic  
253 subtasks first (**Fig. 5A,B**). Curriculum learning has been shown to improve artificial neural net-  
254 work training both in training iterations to convergence and in the final loss (Bengio et al., 2009).  
255 In neuroscience, researchers adopt a wide variety of different curricula to train animals to perform  
256 full experimental tasks. By including curriculum learning, PsychRNN enables researchers to in-  
257 vestigate how training curricula may impact resulting behavioral and neural solutions to cognitive

258 tasks, as well as potentially identify new curricula that may accelerate training. Further, curricula  
259 can be used more broadly to investigate how learning may be influenced and biased by the sets of  
260 tasks an agent has previously encountered.

261 As an example, we trained RNNs on a version of the perceptual decision-making task (from  
262 **Fig. 2**), and examined the effects of using curriculum learning in the training procedure (**Fig.**  
263 **5C,D**). Here, curriculum learning involved initially training the model at high stimulus coherences,  
264 and introducing progressively lower coherences when the model’s performance reached a threshold  
265 level. We found that curriculum learning enabled faster training of models, as commonly observed  
266 in experiments (Krueger and Dayan, 2009).

### 267 **Comparison to other frameworks**

268 PsychRNN compares favorably to alternative high-level frameworks available (**Fig. 6**). Most simi-  
269 lar to PsychRNN is PyCog (Song et al., 2016), another Python package for training RNNs designed  
270 for neuroscientists. PsychRNN presents several key advantages over PyCog. First, PyCog’s back-  
271 end is Theano, which is no longer under active support and development. Second, PyCog has no  
272 native implementation of curriculum learning. Third, task definitions in PyCog are not themselves  
273 modular, making experiments which are trivial to implement in PsychRNN more laborious and  
274 cumbersome for the user. Lastly, PyCog utilizes a built-in “vanilla” stochastic gradient descent  
275 algorithm, whereas PsychRNN allows users to select any optimizer available in the TensorFlow  
276 package.

277 Alternatively research groups may use a general purpose high-level wrapper of TensorFlow, such  
278 as Keras (Chollet and others, 2015), which is not specifically designed for neuroscientific research.  
279 Importantly, these frameworks do not come with any substantial ability to implement biological  
280 constraints. Users interested in testing the impact of such constraints would need to modify the  
281 native Keras Layer objects themselves, which is nontrivial. In addition, Keras does not provide a  
282 framework for modular task definition, which therefore requires the user to translate inputs and  
283 outputs into a form compatible with the model. PsychRNN, by close integration with the Ten-  
284 sorFlow framework, manages to maintain much of the power and flexibility of traditional machine  
285 learning frameworks while also providing custom-built utilities specifically designed for addressing  
286 neuroscientific questions.

### 287 **Discussion**

288 PsychRNN provides a robust and modular package for training RNNs on cognitive tasks, and is  
289 designed to be accessible to researchers with varying levels of deep learning experience. The sepa-  
290 ration into a Python- and NumPy-based **Task** object and a primarily TensorFlow-based Backend  
291 expands access to RNN model training without reducing flexibility and power for users who re-  
292 quire more control over the precise setup of their networks. Further, the modularity of tasks and  
293 network elements enables easy investigation of how task and structure affect learned solution in  
294 RNNs. Lastly, the modular structure facilitates curriculum learning which makes optimization  
295 more efficient and more directly comparable to animal learning.

296 PsychRNN’s modular design enables straightforward implementation of curriculum learning to  
297 facilitate studies of how training trajectories shape network solutions and performance on cogni-  
298 tive tasks. Task shaping is a relatively understudied topic in systems neuroscience, despite its  
299 ubiquity in animal training. For instance, it is poorly understood whether differences in train-  
300 ing trajectories result in different cognitive strategies or neural representations in a task (e.g.,  
301 Latimer and Freedman, 2019). Standardization and automation in animal training may aid exper-  
302 imental investigation of task shaping effects (Murphy et al., 2020; Berger et al., 2018). Although  
303 PsychRNN utilizes a supervised training procedure, rather than reinforcement-based ones used  
304 in animal training, the implementation of curricula enables exploration of how task shaping may  
305 impact learning in cognitive tasks.

306 Future extensions of the PsychRNN codebase can enable investigation into additional neuro-  
307 scientific questions. Some potentially useful directions are the addition of units that exhibit firing  
308 rate adaptation through an internal dynamical variables associated with each unit (Masse et al.,  
309 2019), spiking neurons (Zenke and Ganguli, 2018), and the implementation of networks with short  
310 term associative plasticity (e.g., Miconi, 2017). An interesting area for extending task training  
311 capability is to add trial-by-trial dependencies. In the current version of PsychRNN, each task  
312 trial is trained independently from other trials in the same block. PsychRNN could potentially  
313 be extended to support dependencies across trials by having the loss function and trial specifi-  
314 cation depend on a series of trials. In model training, PsychRNN could be extended to support  
315 learning algorithms apart from supervised gradient descent, such as deep reinforcement learning  
316 algorithms (Botvinick et al., 2020). With the recent release of TensorFlow 2.0 extending function-  
317 ality to match alternative frameworks including PyTorch, we see TF as a strong base on which  
318 to design PsychRNN. PsychRNN allows for future extension to include other frameworks in its  
319 Backend. Importantly, the modular design of PsychRNN can enable such extensions and updates  
320 without forcing any user-side change in task specification or front-end experience.

321 The modular design of PsychRNN also supports extension with various methods for analysis  
322 of trained RNNs, which could be implemented by users. Here, we have provided a basic set of  
323 built-in analysis tools to directly investigate the features and structures of trained RNN weights,  
324 states, and outputs. Because the landscape of analysis methods differs substantially across studies,  
325 built-in analysis methods cannot be comprehensive, and therefore we decided to instead focus on  
326 providing output in forms that are most broadly compatible with common analysis pipelines.

327 The PsychRNN package provides an easy-to-use framework that can be applied and transferred  
328 between research groups to accelerate collaboration and enhance reproducibility. Where in the  
329 current environment research groups need to transfer their entire codebase in order to run an RNN  
330 model, in the PsychRNN framework they are able to transfer just a task or model file for researchers  
331 to investigate and build upon. The ability to test identically specified models across tasks in different  
332 groups, and identically specified tasks across models improves reliability of research. Furthermore,  
333 the many choices in defining and training RNNs can make precise replication of prior published  
334 research difficult. The specification of PsychRNN task files and parameter dictionaries can make  
335 reproduction of RNN studies more open and straightforward.

336 PsychRNN was designed to lower barriers to entry for researchers in neuroscience who are

<sup>337</sup> interested in RNN modeling. In service of this goal we have created a highly user-friendly, clear,  
<sup>338</sup> and modular framework for task specification, while abstracting away much of the deep learning  
<sup>339</sup> background necessary to train and run networks. This modularity also provides access to new  
<sup>340</sup> research directions and a reproducible framework that will aid RNN modeling in neuroscientific  
<sup>341</sup> research.

<sup>342</sup> **Acknowledgements**

<sup>343</sup> This research was supported by NIH grant R01MH112746 (JDM), the Gruber Foundation (DBE),  
<sup>344</sup> and a Barry Goldwater scholarship (JTS). We thank Maxwell Shinn for helpful feedback on the  
<sup>345</sup> package design and manuscript.

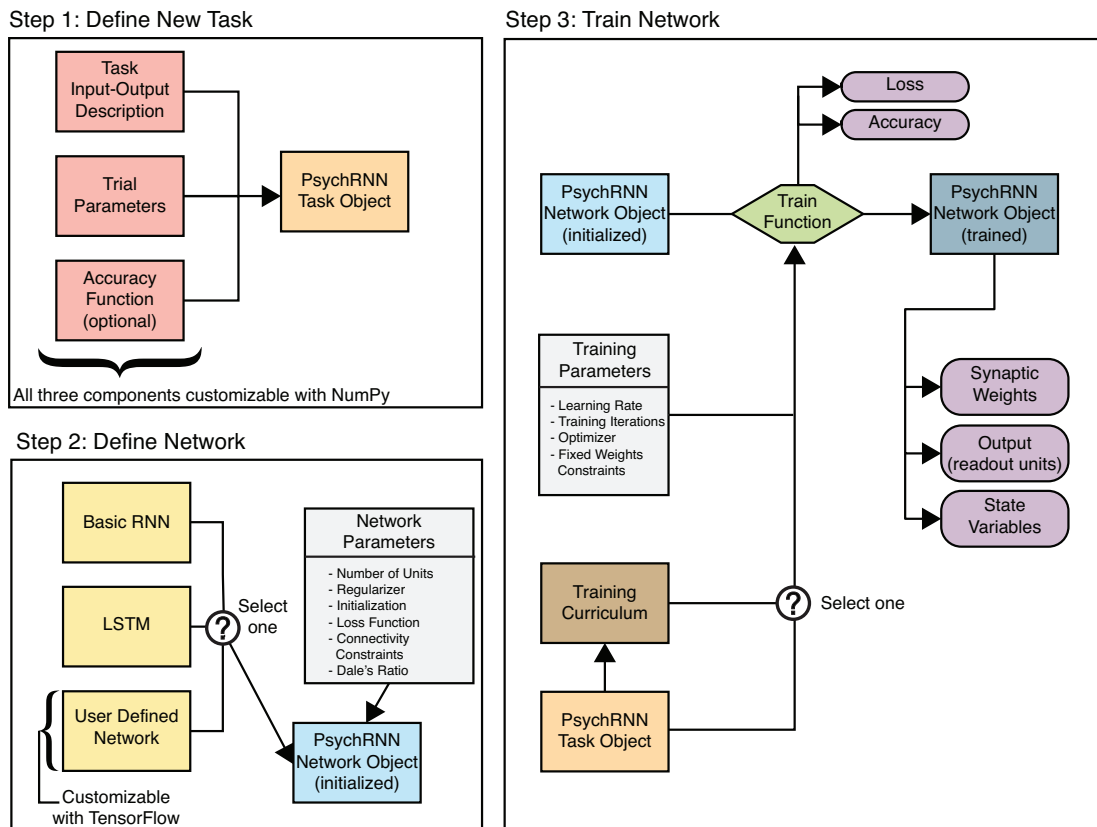
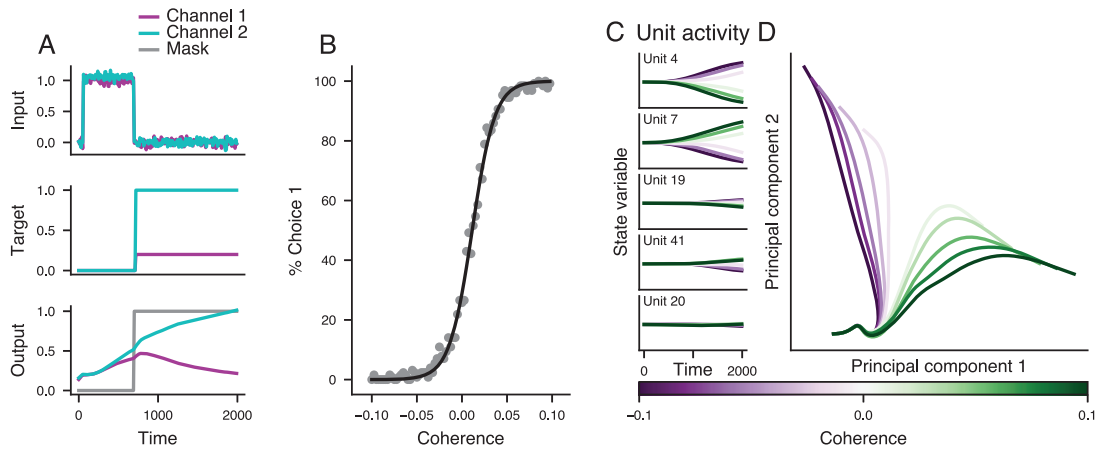


Figure 1: **PsychRNN package structure.** (**Step 1**) Defining a new task requires two NumPy-based components: `trial_function` describes the task inputs and outputs, and `generate_trial_params` defines parameters for a given trial (**Extended Data, Figure 1-1**). Optionally, one can define an accuracy function describing how to calculate whether performance on a trial was successful. (**Step 2**) The Backend defines the network. First, the model, or network architecture, is selected. A basic RNN and LSTM (Hochreiter and Schmidhuber, 1997) are implemented, and more models or architectures can be defined using TensorFlow. That model is then instantiated with a dictionary of parameters, which includes the number of recurrent units and may also include specifications of loss functions, initializations, regularizations, or constraints. If any parameter is not set, a default is used. (**Step 3**) Training specifications, such as the optimizer or curriculum, can be specified. During network training, measures of performance (loss and accuracy) are recorded at regular intervals. Optimization of the network weights is performed to minimize the loss. After training, the synaptic weight matrix can be saved, and state variables and network output can be generated for any given trial.





```

1 from psychrnn.tasks.perceptual_discrimination import PerceptualDiscrimination
2 from psychrnn.backend.models.basic import Basic
3 import tensorflow as tf

4 pdm = PerceptualDiscrimination(dt = 10, tau = 100, T = 2000, N_batch = 128)

5 network_params = pdm.get_task_params() # get the params passed in and defined in pdm
6 network_params['name'] = 'model' # name the model uniquely if running mult models in unison
7 network_params['N_rec'] = 50 # set the number of recurrent units in the model
8 model = Basic(network_params) # instantiate a basic vanilla RNN
9 model.train(pdm) # train model to perform pdm task

10 x,target_output,mask, trial_params = pdm.get_trial_batch() # get pdm task inputs and outputs
11 model_output, model_state = model.test(x) # run the model on input x

```

Figure 2: **Example Task (Perceptual Discrimination)**. (A) Inputs and target output as specified by the task (top two panels), and the network's output for the displayed input (bottom panel). Because the output mask is zero during the stimulus period, the network is not directly constrained during that period. (B) Percent of decisions the network makes for Choice 1 at varying coherence levels. Negative coherence levels indicate stimulus inputs rewarded Choice 2. A psychometric function is fit to the data (black). This plot validates that the network successfully learned the task. (C) State variable activity traces across for a range of stimulus coherences, for multiple example units, averaged over correct trials. The network produces state variable activity across all units. (D) Population activity traces in the subspace of the top two principal components (PCs). PCA was applied to the activity matrix formed by concatenating across coherences the trial-averaged correct-trial traces, for each unit. (E) Minimal example code for using PsychRNN. All relevant modules are imported (lines 1-3), a `PerceptualDiscrimination Task` object is initialized (line 4), the basic RNN model is instantiated and trained (lines 5-9), output and state variables are extracted (lines 10-11).

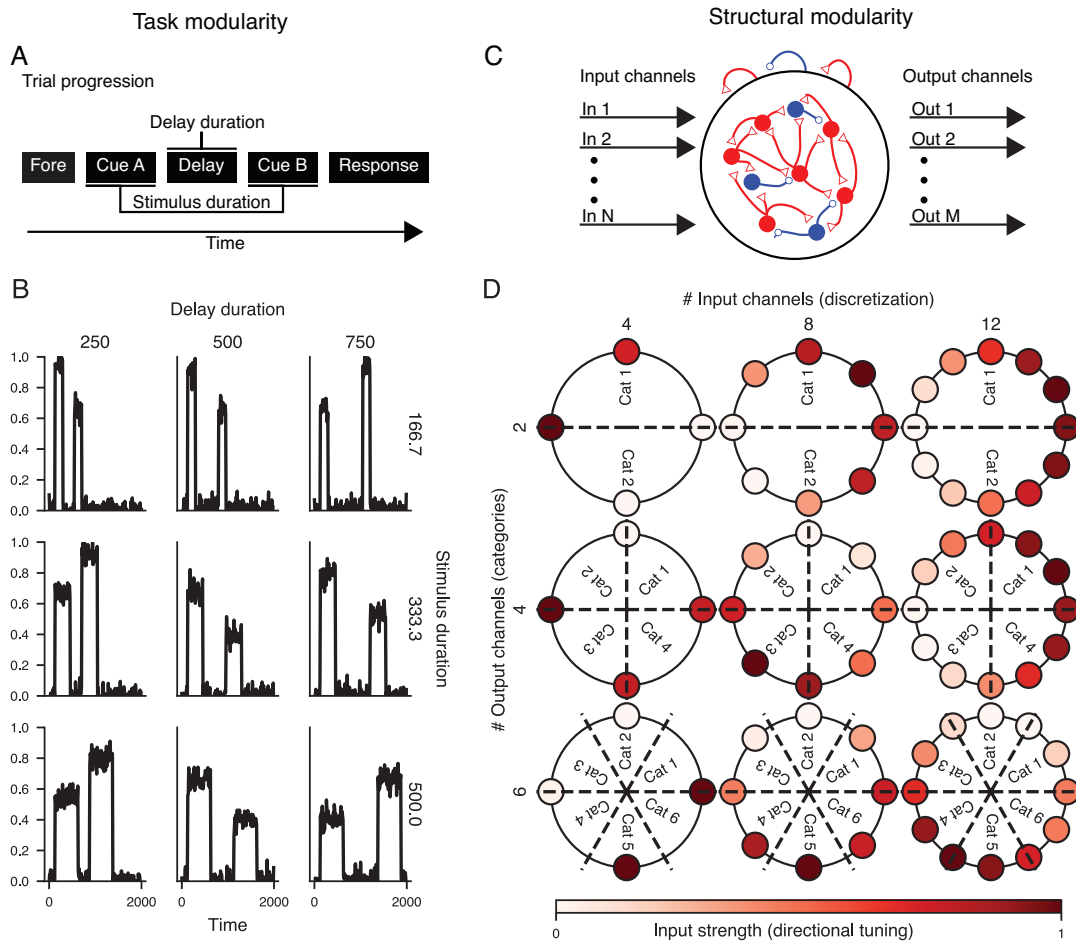


Figure 3: **Modularity of task definition.** (A) Task modularity. This schematic illustrates the trial progression of one trial of a delayed discrimination task. The task is modularly defined such that stimulus and delay duration can be varied easily, simply by changing task parameters. (B) One input channel generated by a delayed discrimination task, with varied stimulus and delay durations (Extended Data, Figure 3-1). Delay duration is varied across columns, and stimulus duration is varied across rows. (C) Structural modularity. Tasks can provide any numbers of channels for input and output on which to train a particular RNN model. Variation in numbers of inputs and outputs is enabled through simple modular task parameters in PsychRNN. (D) Example of a match-to-category task. The number of inputs (colored outer circles) is varied across columns, and the number of output categories (Cat) is varied across rows (Extended Data, Figure 3-2).

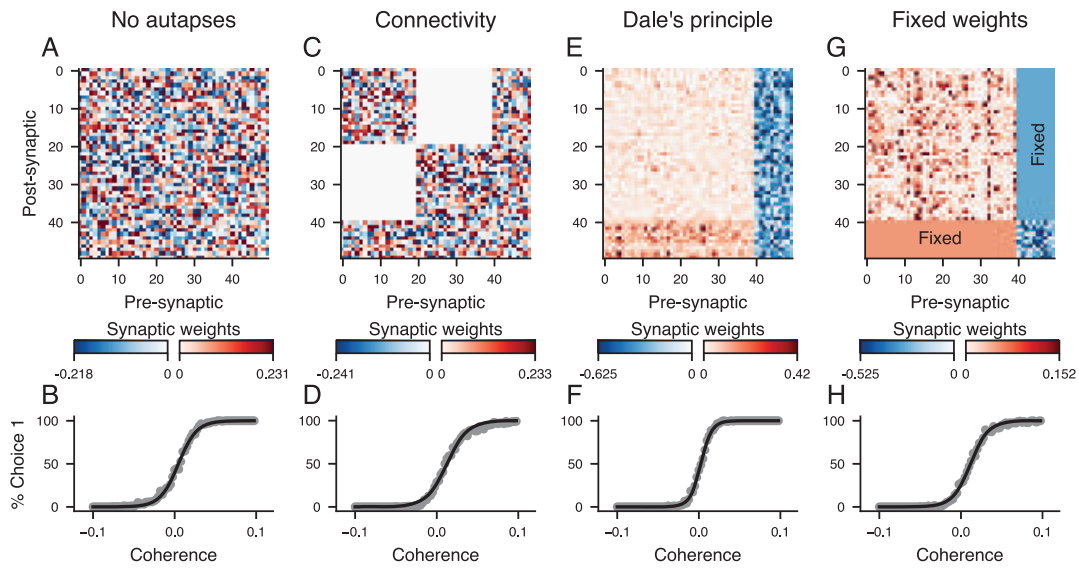


Figure 4: **Neurobiologically motivated constraints.** This figure illustrates the effects of different connectivity constraints on the recurrent weight matrices and psychometric functions of RNNs trained on the perceptual discrimination task (Figure 2). For the recurrent weight matrices (top row), red and blue show excitatory and inhibitory connections, respectively. The coherence plots (bottom row) show that the network successfully trains to perform the task while adhering to the constraints. (A,B) This network is constrained to have no autapses, i.e., no self-connections, as illustrated by zeros along the diagonal of the weight matrix. (C,D) This network is constrained to have two densely connected populations of units with sparse connection between the populations. This constraints can be used to simulate long-range interactions among different brain regions. (E,F) This network is constrained to follow Dale's principle: each neuron either has entirely excitatory or entirely inhibitory outputs. (E,F) This network has Dale's principle enforced and has subset of weights which are fixed, i.e., they cannot be updated by training. In this example, all connections between excitatory and inhibitory neurons are fixed, and other within excitatory-to-excitatory and inhibitory-to-inhibitory connections are plastic during training.

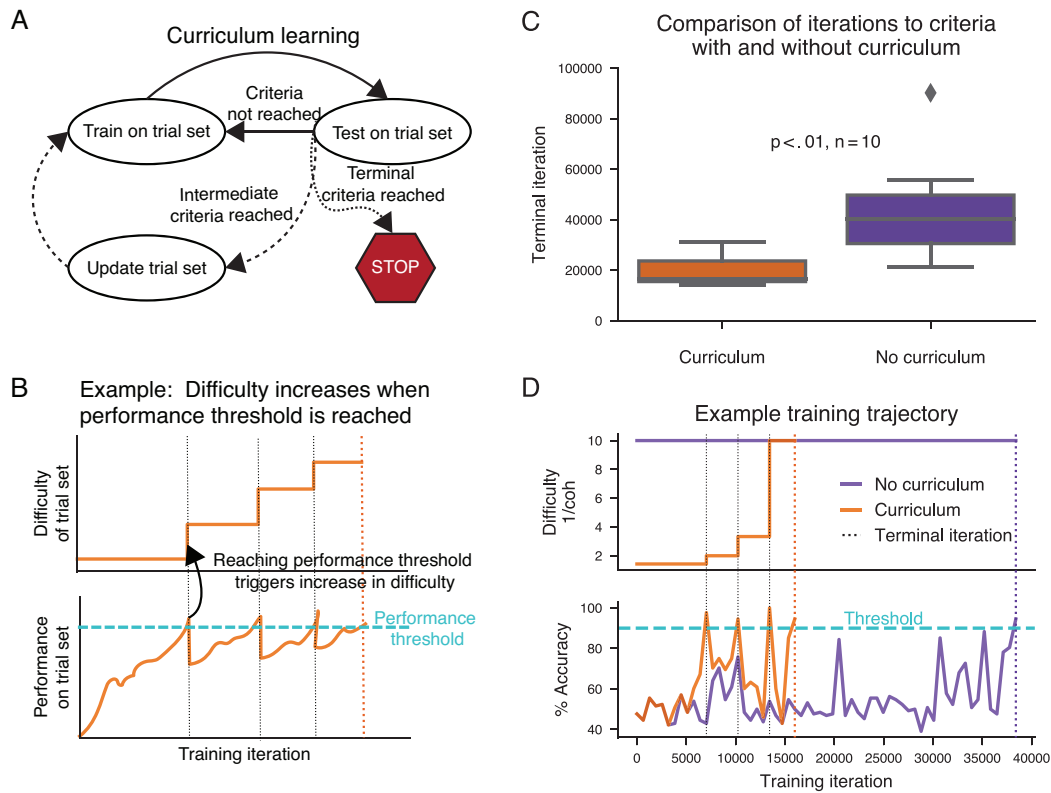


Figure 5: **Curriculum learning.** (A) Schematic of curriculum learning, or task shaping. The network is trained on selections from the trial set, then tested on selections from that trial set. Depending on the performance when testing on the trial set, the trial set can then be updated, e.g., to contain progressively more difficult trial conditions. (B) Example schematic of increasing difficulty of trial set (top) paired with performance over time (bottom). The task difficulty is progressively increased each time performance reaches the performance threshold. (C) Comparison of number of iterations needed to train a network to perform the perceptual discrimination task (from Fig. 2) with 90% accuracy at coherence level of 0.1. Ten networks were randomly initialized and each was trained both on a curriculum with decreasing coherence, and without a curricula with fixed coherence. Networks trained without curriculum learning were trained solely on stimulus with coherence = 0.1. Networks trained with curriculum learning were trained with a curriculum with coherence decreasing from 0.7 to 0.5 to 0.3 to 0.1 as performance improved (see **Extended Data, Figure 5-1**). When the network reached 90% accuracy on stimuli with coherence = 0.1, training was stopped. Networks trained with curriculum learning reached 90% accuracy significantly faster than networks without it ( $p < 0.01$ ). (D) Trajectories of difficulty (defined here as inverse coherence), accuracy, and loss (mean squared error) across training iterations, for two identically initialized networks from (C), one of which was trained with curriculum learning, and one of which was trained without curriculum learning.

Implemented Features	PsychRNN	PyCog (Song et al., 2016)	Keras (Chollet and others, 2015)
Key advantages	Curriculum learning; modular task and network definition; neurobiological constraints; supported backend	Neurobiological constraints on RNNs	Compatible with multiple backends
Language	Python	Python	Python
Backend (actively supported?)	TensorFlow 1 (maintenance mode only) and 2 (yes)	Theano (no)	TensorFlow 1 (maintenance mode only), TensorFlow 2 (yes), Theano (no), CNTK (yes)
Biological constraints supported	Dale's Principle; Connectivity patterns; Fixed-weight training	Dale's Principle; Connectivity patterns; Fixed-weight training	No built-in support
Curriculum learning supported?	Yes	No	Not applicable
Modular task definition?	Yes	No	Not applicable
Object-oriented framework	Task; RNN	RNN	Network layers
LSTM	Built-in support	Not supported	Built-in support
Optimizer	TensorFlow built-in options with implemented regularizers	Stochastic gradient descent (SGD) with implemented regularizers	Built-in options
GPU support	Yes	Yes	Yes
Supports PyTorch	No	No	No

Figure 6: Comparison of PsychRNN to alternative RNN training packages. Red, yellow, and green indicate limited, moderate, and maximal flexibility or accessibility, respectively.

**Extended Data.** PsychRNN package and documentation.

```

1 class SimplePD(Task):
2     def __init__(self, dt, tau, T, N_batch):
3         super(SimplePD, self).__init__(2, 2, dt, tau, T, N_batch)
4     def generate_trial_params(self, batch, trial):
5
6         # -----
7         # Define parameters of a trial
8         # -----
9         params = dict()
10        params['coherence'] = np.random.exponential(scale=1/5)
11        params['direction'] = np.random.choice([0, 1])
12
13        return params
14
15    def trial_function(self, t, params):
16        stim_noise = 0.1
17        onset = self.T/4.0
18        stim_dur = self.T/2.0
19
20        # -----
21        # Initialize with noise
22        # -----
23        x_t = np.sqrt(2*self.alpha*stim_noise*stim_noise)*np.random.randn(self.N_in)
24        y_t = np.zeros(self.N_out)
25        mask_t = np.ones(self.N_out)
26
27        # -----
28        # Retrieve parameters
29        # -----
30        coh = params['coherence']
31        direction = params['direction']
32
33        # -----
34        # Compute values
35        # -----
36        if onset < t < onset + stim_dur:
37            x_t[direction] += 1 + coh
38            x_t[(direction + 1) % 2] += 1
39
40        if t > onset + stim_dur + 20:
41            y_t[direction] = 1.
42
43        if t < onset + stim_dur:
44            mask_t = np.zeros(self.N_out)
45
46        return x_t, y_t, mask_t

```

**Extended Data, Figure 1-1: Example PsychRNN code showing task definition.** This code sample defines a simple two-alternative perceptual decision-making task. The function `generate_trial_params` selects the coherence and direction on a trial by trial basis. `trial_function` sets the input, target output and output mask depending on the time in the trial and the passed parameters.

```
1 from psychrnn.tasks.delayed_discrim import DelayedDiscrimination
2
3 for i in range(3):
4     for j in range(3):
5         dd = DelayedDiscrimination(dt = 10, # simulation time step
6                                   tau = 100, # unit time constant
7                                   T = 2000, # trial length
8                                   N_batch = 1, # number of trials per update
9                                   delay_duration = (j+1) * 250, # delay length
10                                  decision_duration = 250, # decision length
11                                  onset_time = 125, # first stimulus onset time
12                                  stim_duration_1 = (i+1)/3 * 500, # stim 1 length
13                                  stim_duration_2 = (i+1)/3 * 500) # stim 2 length
14     x, target_output, mask, trial_params= dd.get_trial_batch() # get task
```

**Extended Data, Figure 3-1: Example PsychRNN code showing modularity of task structure.** This code sample produces all data shown in **Fig. 3B**. Although each plot in **Fig. 3B** has different task durations, iteration through all of them is simplified through the object-oriented modular task definitions enabled by PsychRNN. Resulting code is compact and readable compared to non-modular alternatives.



```
1 from psychrnn.tasks.match_to_category import MatchToCategory
2
3 for i in range(3):
4     for j in range(3):
5         mc = MatchToCategory(dt=10, # simulation time step
6                               tau = 100, # unit time constant
7                               T = 2000, # trial length
8                               N_batch = 1, # number of trials per training update
9                               N_in=(i+1)*4, # number of network inputs
10                              N_out =(j +1)*2) # number of network outputs
11         x, target_output, mask, trial_params = mc.get_trial_batch() # get task
```

**Extended Data, Figure 3-2: Example PsychRNN code showing modularity of task inputs and outputs.** This code sample produces all data shown in Figure 3D. Although each plot in **Fig. 3D** has different numbers of inputs and outputs, iteration through all of them is simplified through the object-oriented modular task definitions enabled by PsychRNN. Resulting code is compact and readable compared to non-modular alternatives.

```

1 from psychrnn.tasks.perceptual_discrimination import PerceptualDiscrimination
2 from psychrnn.Backend.curriculum import Curriculum
3
4 coherences = [.7, .5, .3, .1]
5 task_list = [PerceptualDiscrimination(dt=10, # simulation time step
6     tau = 100, # unit time constant
7     T = 2000, # trial length
8     N_batch = 128, # number of trials per update
9     coherence = coh # coherence of trials
10    ) for coh in coherences]
11 curriculum = Curriculum(task_list, # list of tasks that make up the curriculum
12    output_file="./accuracies", # path to save metric value
13    metric_epoch= 5 # interval to check the metric
14    )
15 train_params = {
16     "save_weights_path": "./weights.npz", # path to save trained network weights
17     "training_iters": 100000, # maximum number of training iterations
18     "loss_epoch": 5, # how often to calculate loss
19     "curriculum": curriculum # Curriculum object
20 }
21
22 # curriculum-training specific wrapper of the train function
23 losses , training_time, initialization_time = model.train_curric(
24     train_params # training parameters
25 )

```

**Extended Data, Figure 5-1: Example PsychRNN code showing curriculum learning.**

This code sample trains an RNN on a sequence of perceptual discrimination tasks with decreasing stimulus coherence. The network is first trained to perform the task with high coherence. Once the network reaches 90% accuracy on a given task (here, set of stimulus coherences), the network initiates training on the next task. This continues until the network has reached 90% accuracy on the final task—in this case, the lowest-coherence task. Curriculum learning is implemented by defining a list of tasks that form the curriculum (line 4-10), and passing that list in to the `Curriculum` class to form a `Curriculum` object (line 11-14). That `Curriculum` object is then included in the training parameters dictionary (line 15-20), and when the network is passed those training parameters for training, the network will be trained using the curriculum, or sequence of task parameters defined in lines 4-10.

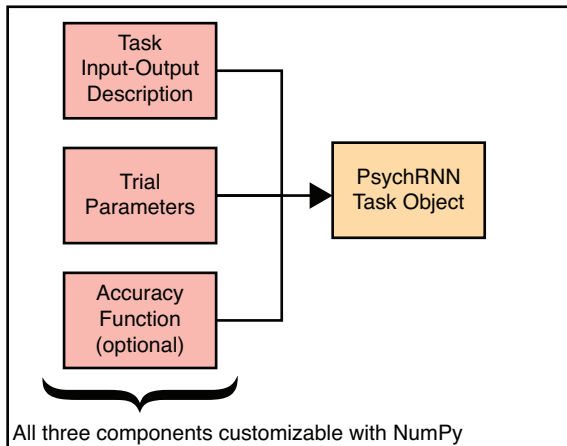
346 **References**

- 347 Barak O (2017) Recurrent neural networks as versatile tools of neuroscience research. *Curr Opin*  
348 *Neurobiol* 46:1–6.
- 349 Bengio Y, Louradour J, Collobert R, Weston J (2009) Curriculum learning In *Proceedings of the*  
350 *26th Annual International Conference on Machine Learning, ICML '09*, pp. 41–48, New York,  
351 NY, USA. ACM.
- 352 Berger M, Calapai A, Stephan V, Niessing M, Burchardt L, Gail A, Treue S (2018) Standardized  
353 automated training of rhesus monkeys for neuroscience research in their housing environment. *J*  
354 *Neurophysiol* 119:796–807.
- 355 Botvinick M, Wang JX, Dabney W, Miller KJ, Kurth-Nelson Z (2020) Deep reinforcement learning  
356 and its neuroscientific implications. *Neuron* 107:603–616.
- 357 Carnevale F, de Lafuente V, Romo R, Barak O, Parga N (2015) Dynamic control of response  
358 criterion in premotor cortex during perceptual detection under temporal uncertainty. *Neu-*  
359 *ron* 86:1067–1077.
- 360 Chollet F et al. (2015) Keras <https://github.com/fchollet/keras>.
- 361 Freedman DJ, Assad JA (2006) Experience-dependent representation of visual categories in parietal  
362 cortex. *Nature* 443:85–88.
- 363 Glorot X, Bengio Y (2010) Understanding the difficulty of training deep feedforward neural networks  
364 In Teh YW, Titterton M, editors, *Proceedings of the Thirteenth International Conference*  
365 *on Artificial Intelligence and Statistics*, Vol. 9 of *Proceedings of Machine Learning Research*,  
366 pp. 249–256, Chia Laguna Resort, Sardinia, Italy. PMLR.
- 367 Hochreiter S, Schmidhuber J (1997) Long short-term memory. *Neural Computation* 9:1735–1780.
- 368 Kriegeskorte N (2015) Deep neural networks: A new framework for modeling biological vision and  
369 brain information processing. *Annu Rev Vis Sci* 1:417–446.
- 370 Krueger KA, Dayan P (2009) Flexible shaping: how learning in small steps helps. *Cogni-*  
371 *tion* 110:380–94.
- 372 Latimer KW, Freedman DJ (2019) Learning dependency of motion direction tuning in the lateral  
373 intraparietal area during a categorization task. Program No. 756.10. 2019 Neuroscience Meeting  
374 Planner. Chicago, IL: Society for Neuroscience, 2019. Online.
- 375 Mante V, Sussillo D, Shenoy KV, Newsome WT (2013) Context-dependent computation by recur-  
376 rent dynamics in prefrontal cortex. *Nature* 503:78–84.
- 377 Masse NY, Yang GR, Song HF, Wang XJ, Freedman DJ (2019) Circuit mechanisms for the  
378 maintenance and manipulation of information in working memory. *Nat Neurosci* 22:1159–1167.

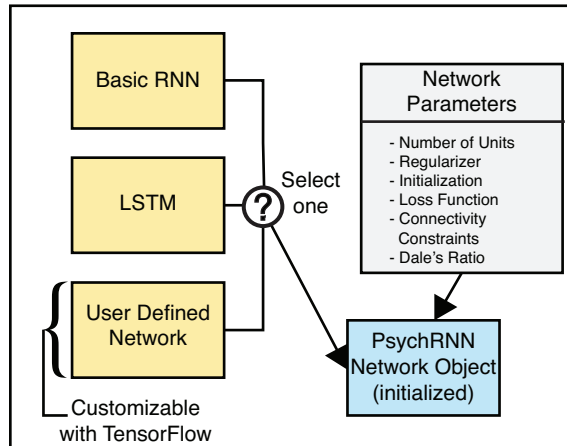
- 379 Miconi T (2017) Biologically plausible learning in recurrent neural networks reproduces neural  
380 dynamics observed during cognitive tasks. *Elife* 6.
- 381 Murphy TH, Michelson NJ, Boyd JD, Fong T, Bolanos LA, Bierbrauer D, Siu T, Balbi M, Bolanos  
382 F, Vanni M, LeDue JM (2020) Automated task training and longitudinal monitoring of mouse  
383 mesoscale cortical circuits using home cages. *Elife* 9.
- 384 Orhan AE, Ma WJ (2019) A diverse range of factors affect the nature of neural representations  
385 underlying short-term memory. *Nat Neurosci* 22:275–283.
- 386 Rajan K, Harvey CD, Tank DW (2016) Recurrent network models of sequence generation and  
387 memory. *Neuron* 90:128–42.
- 388 Remington ED, Narain D, Hosseini EA, Jazayeri M (2018) Flexible sensorimotor computations  
389 through rapid reconfiguration of cortical dynamics. *Neuron* 98:1005–1019.e5.
- 390 Richards BA, Lillicrap TP, Beaudoin P, Bengio Y, Bogacz R, Christensen A, Clopath C, Costa  
391 RP, de Berker A, Ganguli S, Gillon CJ, Hafner D, Kepecs A, Kriegeskorte N, Latham P, Lindsay  
392 GW, Miller KD, Naud R, Pack CC, Poirazi P, Roelfsema P, Sacramento J, Saxe A, Scellier B,  
393 Schapiro AC, Senn W, Wayne G, Yamins D, Zenke F, Zylberberg J, Therien D, Kording KP  
394 (2019) A deep learning framework for neuroscience. *Nat Neurosci* 22:1761–1770.
- 395 Rikhye RV, Gilra A, Halassa MM (2018) Thalamic regulation of switching between cortical repre-  
396 sentations enables cognitive flexibility. *Nat Neurosci* 21:1753–1763.
- 397 Roitman JD, Shadlen MN (2002) Response of neurons in the lateral intraparietal area during a  
398 combined visual discrimination reaction time task. *J Neurosci* 22:9475–89.
- 399 Romo R, Brody CD, Hernández A, Lemus L (1999) Neuronal correlates of parametric working  
400 memory in the prefrontal cortex. *Nature* 399:470–3.
- 401 Ruder S (2017) An overview of gradient descent optimization algorithms. *arXiv* p. 1609.04747.
- 402 Song HF, Yang GR, Wang XJ (2016) Training excitatory-inhibitory recurrent neural networks for  
403 cognitive tasks: A simple and flexible framework. *PLoS Comput Biol* 12:e1004792.
- 404 Sussillo D (2014) Neural circuits as computational dynamical systems. *Curr Opin Neuro-*  
405 *biol* 25:156–63.
- 406 Sussillo D, Abbott L (2009) Generating coherent patterns of activity from chaotic neural networks.  
407 *Neuron* 63:544 – 557.
- 408 Sussillo D, Churchland MM, Kaufman MT, Shenoy KV (2015) A neural network that finds a  
409 naturalistic solution for the production of muscle activity. *Nature neuroscience* 18:1025.
- 410 V. Le Q, Jaitly N, E. Hinton G (2015) A simple way to initialize recurrent networks of rectified  
411 linear units .

- 412 Yamins DLK, DiCarlo JJ (2016) Using goal-driven deep learning models to understand sensory  
413 cortex. *Nat Neurosci* 19:356–65.
- 414 Yang GR, Wang XJ (2020) Artificial neural networks for neuroscientists: A primer. *Neu-*  
415 *ron* 107:1048–1070.
- 416 Zenke F, Ganguli S (2018) Superspike: Supervised learning in multilayer spiking neural networks.  
417 *Neural Comput* 30:1514–1541.

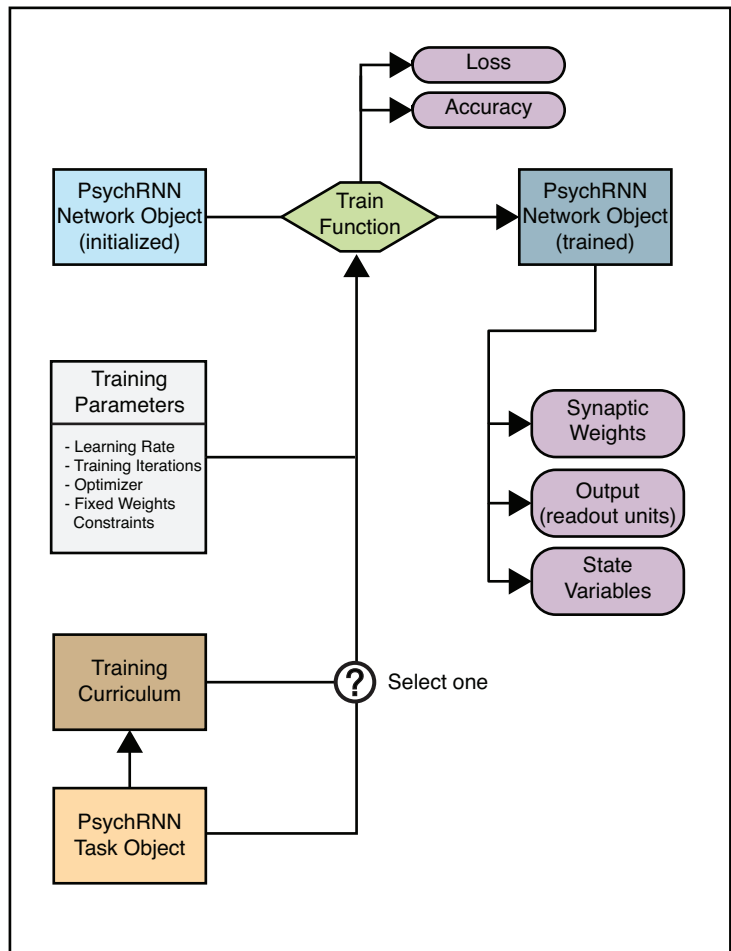
Step 1: Define New Task

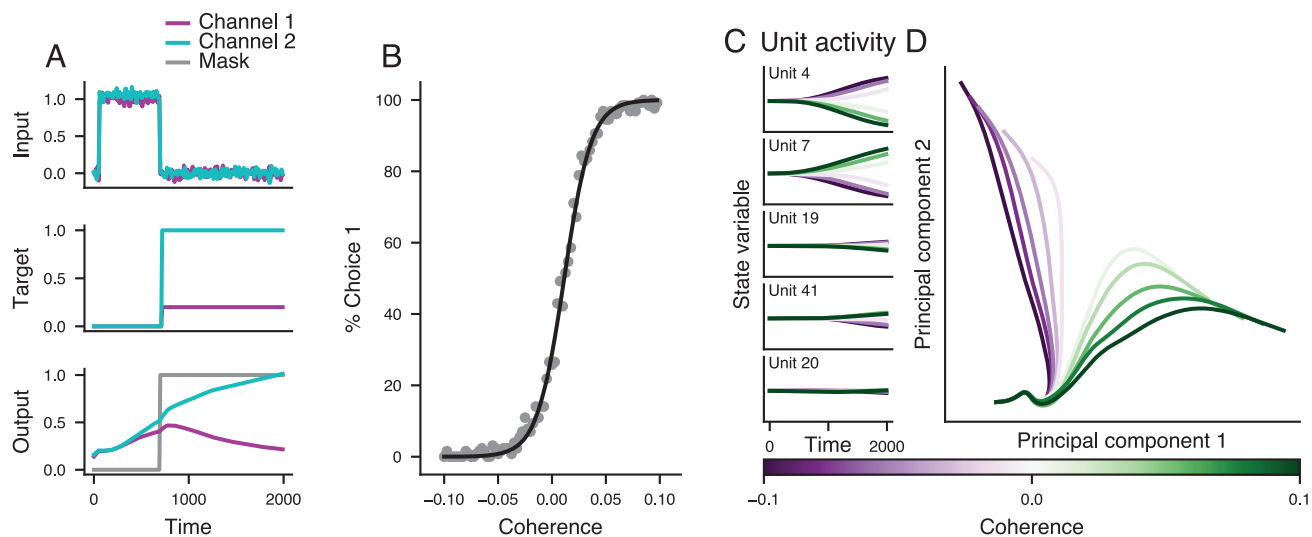


Step 2: Define Network



Step 3: Train Network





E

```




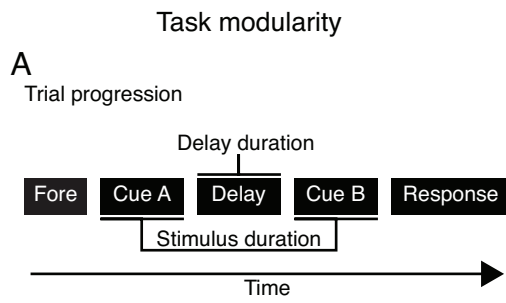
1 from psychrnn.tasks.perceptual_discrimination import PerceptualDiscrimination
2 from psychrnn.backend.models.basic import Basic
3 import tensorflow as tf

4 pdm = PerceptualDiscrimination(dt = 10, tau = 100, T = 2000, N_batch = 128)

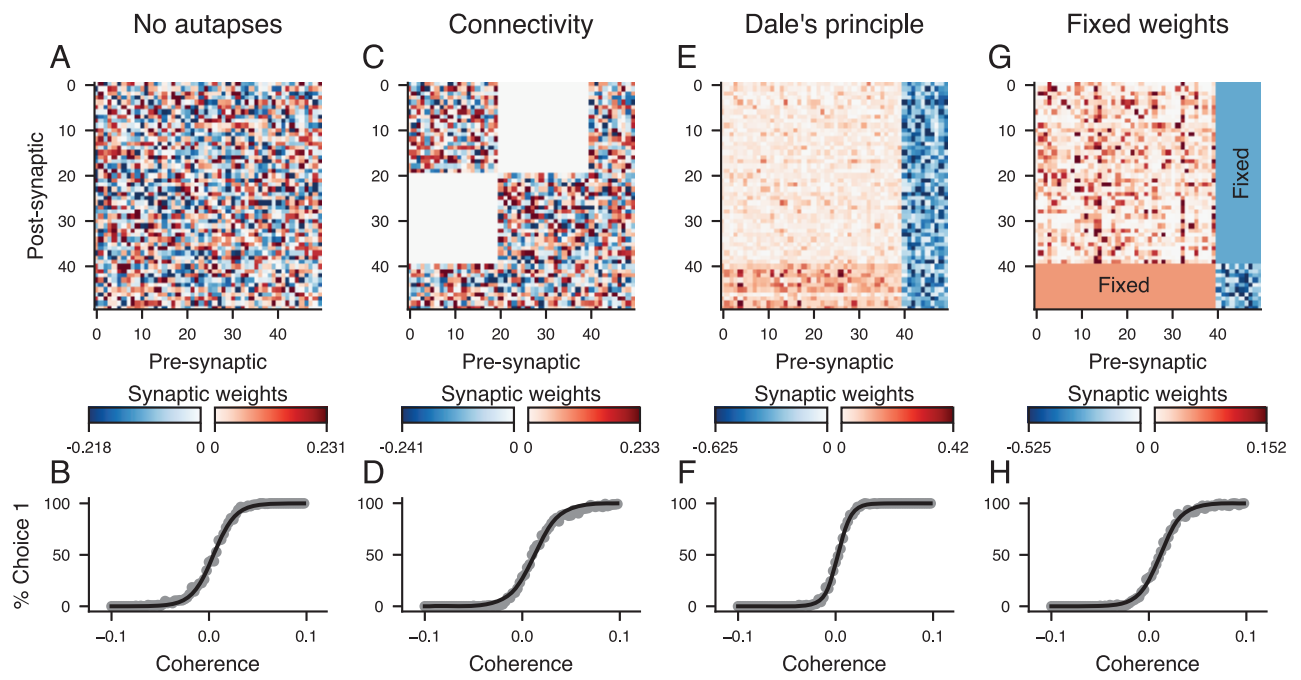
5 network_params = pdm.get_task_params() # get the params passed in and defined in pdm
6 network_params['name'] = 'model' # name the model uniquely if running mult models in unison
7 network_params['N_rec'] = 50 # set the number of recurrent units in the model
8 model = Basic(network_params) # instantiate a basic vanilla RNN
9 model.train(pdm) # train model to perform pdm task

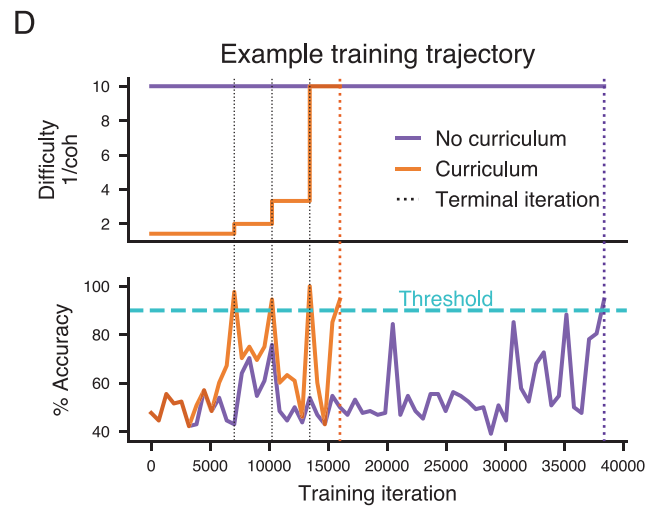
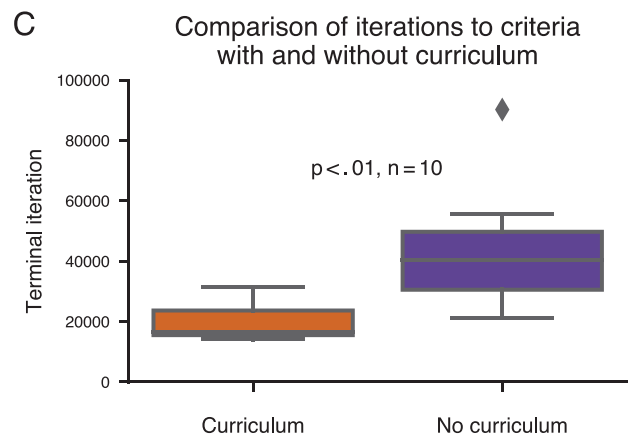
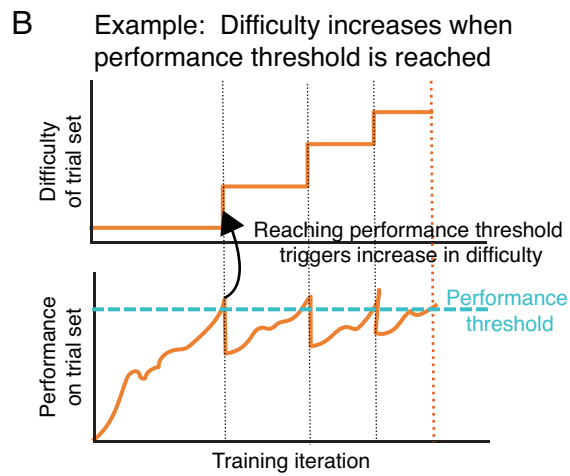
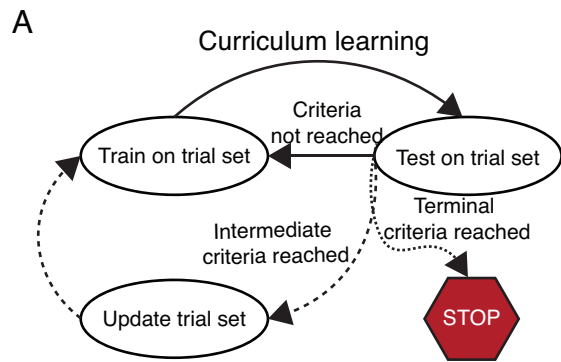
10 x,target_output,mask, trial_params = pdm.get_trial_batch() # get pdm task inputs and outputs
11 model_output, model_state = model.test(x) # run the model on input x

```









Implemented Features	PsychRNN	PyCog (Song et al., 2016)	Keras (Chollet and others, 2015)
Key advantages	Curriculum learning; modular task and network definition; neurobiological constraints; supported backend	Neurobiological constraints on RNNs	Compatible with multiple backends
Language	Python	Python	Python
Backend (actively supported?)	TensorFlow 1 (maintenance mode only) and 2 (yes)	Theano (no)	TensorFlow 1 (maintenance mode only), TensorFlow 2 (yes), Theano (no), CNTK (yes)
Biological constraints supported	Dale's Principle; Connectivity patterns; Fixed-weight training	Dale's Principle; Connectivity patterns; Fixed-weight training	No built-in support
Curriculum learning supported?	Yes	No	Not applicable
Modular task definition?	Yes	No	Not applicable
Object-oriented framework	Task; RNN	RNN	Network layers
LSTM	Built-in support	Not supported	Built-in support
Optimizer	TensorFlow built-in options with implemented regularizers	Stochastic gradient descent (SGD) with implemented regularizers	Built-in options
GPU support	Yes	Yes	Yes
Supports PyTorch	No	No	No